

Real-time Hazard Detection for Landers

Andres Huertas, Yang Cheng and Larry H. Matthies
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

Abstract— Unmanned planetary landers to date have landed "blind"; that is, without the benefit of onboard landing hazard detection and avoidance systems. This constrains landing site selection to very benign terrain, which in turn constrains the scientific agenda of missions. The state of the art Entry, Descent, and Landing (EDL) technology can land a spacecraft on Mars somewhere within a very large landing ellipse (20-100 km). However, even if a landing ellipse is only a few kilometers long, it is very likely to contain hazards such as craters, discontinuities, steep slopes, and large rocks, regardless of how the ellipse is selected. A lander that encounters a large rock, falls off a cliff, or tips over on a steep slope can sustain mission-fatal damage. In this paper, we will briefly review sensor options for landing hazard detection and identify an approach based on stereo vision and shadow analysis that addresses the broadest set of missions. We present the hazard detection approach which fuses stereo vision and shadow based rock detection to maximize the spacecraft safety. We discuss in detail several performance models for slope estimation and rock detection within this approach and validate those models experimentally. Instantiating our model of rock detection reliability for Mars predicts that this approach will reduce the probability of failed landing by at least a factor of 4 in any given terrain. We will also discuss a recent activity to convert the shadow-based rock detector into a rock detection and mapping tool that we are using to process very large, high-resolution HiRISE images from the Mars Reconnaissance Orbiter (MRO) to assist in landing site selection for the Phoenix mission.

I. INTRODUCTION

Landing site selection procedures in planetary exploration use all available remote sensing data to characterize the safety of potential sites before landing is attempted. With cameras now in orbit around Mars and planned to orbit Earth's Moon, it is possible to map all landing hazards larger than a few meters across. Planned precision navigation capabilities will allow avoiding such hazards based only on orbital mapping. However, slopes on the scale of a lander (e.g. < 6 m across) and rocks that could be fatal to a lander (eg. < 3 m in diameter and > 50 cm tall) may not be detected from orbit. Many sites of scientific interest on Mars, in the lunar highlands, and on other moons and asteroids have rock distributions high enough to create a landing failure probability of several percent for blind landers. In contrast, the Mars Science Laboratory (MSL) lander/rover in development for a 2009 launch will accept a landing failure probability due to rock impalement of only 0.25%. For a blind landing, this rules out well over half the surface of the planet. Recent imaging from the HiRISE camera on the Mars Reconnaissance Orbiter (MRO) for Phoenix mission landing site selection revealed high boulder concentrations near the

North Pole of Mars, areas previously considered benign for a lander. Therefore, increasing the accessible surface area requires even higher resolution orbital imagery and/or onboard landing hazard detection (HD) and avoidance capabilities.

Sensors options for HD have been studied for many years, including lidar, radar, and passive imaging [1,2]. Lidar and radar are attractive because they are direct ranging sensors applicable at relatively high altitudes. However, many factors make passive imaging attractive, including a shorter development cycle, potential for smaller size, lower power consumption and lower cost [2]. Landers typically carry descent cameras for scientific imaging that could also be used for HD. A navigation camera may also be needed at high altitude for landmark recognition for precision navigation. Such camera can also be used for HD.

There are still many passive imaging options, including use of color, texture, shading, structure from motion (SFM), stereo, and visible vs. thermal spectral bands. Any selected option also must have a statistical model of hazard detection performance that has been validated experimentally. The goal of modeling is to show that the probability of landing failure is within acceptable limits.

Section II examines planetary landing scenarios to identify a set of sensor/algorithm alternatives with broadest applicability and to determine nominal sensor performance requirements. The conclusion is that stereo vision and shadow analysis appear to cover the widest set of missions with the least complexity. Section III summarizes algorithms we have developed to date for slope estimation and rock detection with these sensing modalities. Sections IV and V describe performance modeling and evaluation work for stereo-based and shadow-based hazard detection. Section V incorporates these results into an overall model of safe landing probability with these sensors. This work also represents a case study in vision system reliability modeling for autonomous navigation that is applicable to lidar and may be valuable in other contexts.

II. LANDING SCENARIOS AND SENSOR OPTIONS

Mars is one of the most challenging places to do landing hazard detection because the rapid descent affords a short time for hazard detection and because the atmosphere constraints when sensing can be done and also reduces image contrast. Thus we use a Mars landing scenario as a design driver, since solutions that work from Mars should apply to most target bodies.

The descent sequence designed for the upcoming MSL mission provides a well-defined reference scenario. This includes a lateral divert maneuver starting about 1.2 km above ground level (AGL) and ending about 100 m AGL to get clear of the parachute; the lateral movement covers about 25% of the starting altitude. Doing precise terrain relative navigation (TRN) by map matching before this point will allow such a maneuver to be targeted to avoid large hazards known from orbital reconnaissance, such as craters up to ~ 100 -200 m in diameter [3]. Detecting small scale hazards before or during this maneuver is impractical for several reasons: (1) it would be expensive because it would require very high sensor angular resolution over a wide field of regard, (2) it would require very accurate navigation to guarantee avoiding all small scale hazards from more than 1 km away, and (3) during the maneuver the high spacecraft attitude rates would make it difficult to obtain low smear, high SNR terrain images aimed at the right place(s) on the ground. At the end of this divert, descent is vertical and relatively slow, so HD is possible at this point to enable a second maneuver of 1-2 lander diameters to avoid small-scale hazards, such as rocks. Thus, performing HD at or below ~ 100 m AGL appears to be most practical for MSL-like missions.

With descent imagery, color, texture, and shape from shading are not promising for HD for a variety of reasons, including results from prior missions that show negligible color variation on asteroid Eros [4] and the impracticality of getting metric slope and rock size information with sufficient accuracy from texture and shading. Contrast in thermal imagery can discriminate rocks from soil over part of the diurnal cycle [5]. However, to minimize cost we would like HD and landmark matching to use the same camera; since the vast majority and the highest resolution orbital mapping imagery is visible spectrum, this is a disadvantage to using thermal imagery for HD.

Shadows can be used to recognize hazardous rocks from altitudes of 1 km or more [2], but this does not enable slope estimation. SFM can enable slope and rock detection if maneuvers are practical that give adequate parallax and enable aiming the camera at the landing site from two or more locations on the descent trajectory. This may be practical for missions to small bodies, like comets and asteroids, but it is costly and difficult for large bodies, like Mars. Binocular stereo baselines of ~ 1 m or more appear to be feasible for most landers and can enable slope and rock detection at altitudes up to about 100 m. Given that this fits the challenging reference mission scenario described above, stereo vision is our primary approach. Shadow analysis can augment rock detection for small incremental runtime cost and can significantly increase rock detection altitude for missions where that is needed, so we include shadows in our approach. Based on our current knowledge of hazard densities around the solar system, this approach is applicable to most or all lander missions. As we discuss below, the speed, reliability, and hardware maturity of this approach makes it a candidate for missions in about five years.

Interest remains in lidar for HD, particularly for robotic landers in permanently dark regions of the lunar poles and for crewed landers; however, it appears lidar is further from

maturity for lander applications. The HD algorithms and performance modeling we apply to range data from stereo are applicable to lidar as well.

III. SUMMARY OF HAZARD DETECTION ALGORITHMS

We have developed three vision algorithms for the small scale hazard detection: (1) Stereo-based slope estimation; (2) shadow-based rock detection and (3) stereo-based rock detection. This section briefly summarizes the algorithms; the following section describes their performance.

Fig. 1 illustrates the dataset (wall dataset) used in most of our experiments on hazard detection. The stereo rig included two 1600x1200 cameras with a 1 m baseline. The cameras horizontal and vertical FOVs are 22° and 18° respectively. At least 30 images were collected every 10m from 10m to 100m “altitude.” Ground truth range data was collected using a Leica Total Station. We also recorded Sun azimuth and elevation angles to supplement evaluation of shadow-based rock detection.

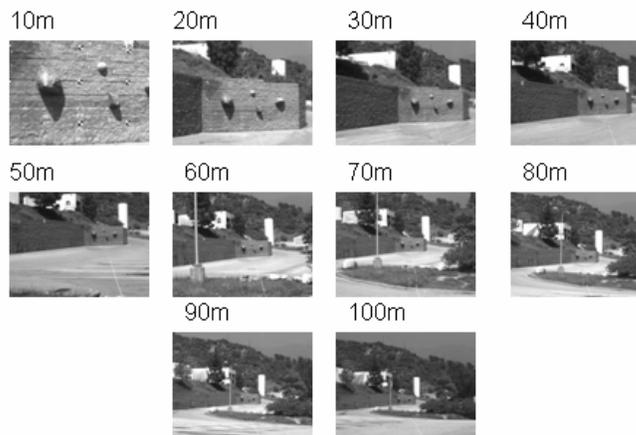


Fig. 1. The wall dataset with ground truth used to evaluate stereo-based slope and rock detection and to supplement evaluation of shadow-based rock detection.

A. Stereo-based Slope Estimation

We use a real-time stereo algorithm that uses five overlapping correlation windows (SAD5) to improve range data quality along object borders [6]. We are currently implementing this algorithm in field programmable gate arrays (FPGAs) and expect to be able to make it operate on 1024x1024 pixel imagery at 10 frames/second (fps) or more [7]. Fig. 2 illustrates a stereo result applied to the wall data at 40 m altitude.

The slope estimation algorithm uses the range data from stereo matching to produce a slope estimate by robust plane fitting. The algorithm has been tested with data that simulates “altitudes” up to 100m to produce slope error vs. latitude assessments relative to lander scale slopes. The algorithm consists of two steps:

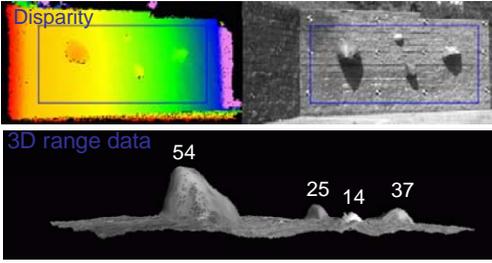


Fig. 2. Sample SAD5 stereo vision range imaging results. Upper right: brick wall with several synthetic rocks viewed from 40 m distance. Upper left: false color range image; red is closest and magenta is furthest. The overlaid rectangle shows the area used to evaluate plane fitting for slope estimation. Bottom: 3-D rendering from below. The numbers above the rocks denote their true height in cm.

The slope estimation algorithm uses the range data from stereo matching to produce a slope estimate by robust plane fitting. The algorithm has been tested with data that simulates “altitudes” up to 100m to produce slope error vs. latitude assessments relative to lander scale slopes. The algorithm consists of two steps:

- 1) An elevation map is produced by SAD5 stereo matching.
- 2) A slope estimate is obtained from robust plane fitting. For this we first perform a least median square fit that includes the rocks (outliers) on the surface. The process is repeated for multiple triplets of points. If the median of the squared plane error is a minimum, we keep these points. Next we discard points far from the plane and apply a least squared fit to the remaining points to obtain the slope estimate. Fig. 3 shows an example of plane fitting (red) applied to the **wall** range data (white) at 30m and at 70m.

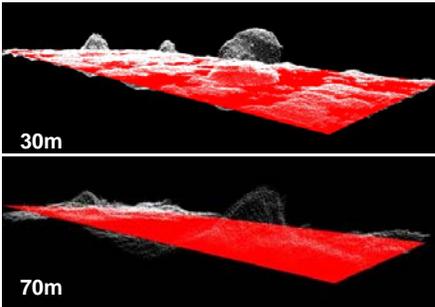


Fig. 3. Underlying surface plane fit applied to 3D range data from SAD5 stereo for two different altitudes.

B. Stereo-based Rock Detection

We first apply SAD5 stereo to obtain an elevation map and underlying surface plane. Then we apply rock detection in four steps (see Fig. 4):

- 1) Threshold the residuals from a robust plane fit.
- 2) The regions over the 1σ threshold are extracted. Then we estimate the deviation from the plane fit.
- 3) Extract potential rocks from connected components and discard noise regions.

- 4) Estimate rock height and position by averaging the n highest range points in each region. We typically use $n = 25$. This reduces noise in the estimates considerably.

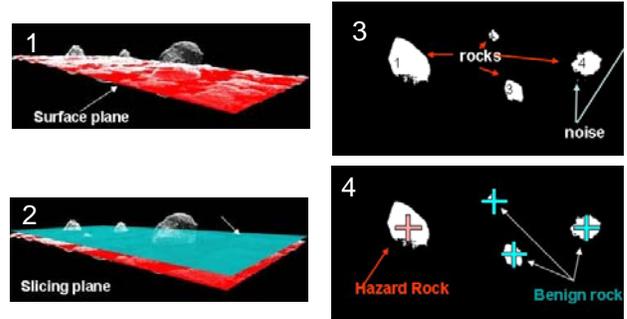


Fig. 4. Stereo-based rock detection: 1) Robust plane fit to range data. 2) Threshold fit residuals at 1σ above the surface plane. 3) Extract potential connected rock points. 4) Remove noise (small) regions and estimate rock height and position.

C. Shadow-based Rock Detection

The ability to detect rock hazards at much higher altitudes than stereo can enable early detection of rocks in the scale of the lander and thus enables early assessment of hazards and diverts operations with reduced effects on the fuel budget. A possible target mega-pixel sensor (12° FOV) images a 1m diameter rock in 5 pixels at 1000 m, enables such early for rock hazard detection. The cues to the presence of rocks are the shadows they cast. Although restricted in applicability by a suitable range of illumination angles, by available shadow saliency, and by the imager resolution, the expectations for these combine favorably to enable this relatively straightforward vision task to become useful to landers in general.

The shadow-based rock detection algorithm has been described in detail in [2] and we only summarize it here. It consists of four steps, illustrated in Fig. 5:

- 1) Image acquisition and state instantiation. State instantiation refers to an environmental state (Sun angles, shadow contrast) and a spacecraft state (altitude, pose).
- 2) Shadow Segmentation. Shadow regions are segmented by applying a modified Maximum Entropy Thresholding (gMET) algorithm [2]. In this method, the shadow regions are segmented by analysis of the histogram of a modified version of the input image. The modified image is computed by adding a gamma-corrected version of the image to the original image and allowing bright areas (non-shadows) to saturate at the image bit-depth.
- 3) Shadow Analysis. The aim here is not to derive accurately delineated rock models thus we do not perform explicit shadow casting to shadow cast correspondences. We also expect some measure of shadow blending and merging, and we do not make attempts to segment the shadows individually. Instead, we fit a “best-ellipse” as described in [2] to the shadow regions that are larger than a certain size (typically 5 pixels) and process those. The local ground sampling distance (GSD) is known from

altitude measurements and sensor specs.

4) Rock Modeling. A circular cross-section model (cyan in Fig. 5) is sufficient for our purposes. The parameters of the shadow elliptical approximation are combined with the Sun angle information to estimate shadow length and width, and rock height and location.

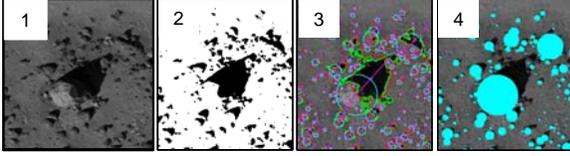


Fig. 5. Sample results for shadow-based rock detection. 1) The input image represents the data and associated environmental conditions, lander parameters, spacecraft state. 2) Shadow segmentation. 3) Shadow analysis, 4) Rock model (position, diameter and height).

IV. HAZARD DETECTION PERFORMANCE

A. Slope Error Analysis

The slope uncertainty model incorporates the following six factors (see Fig. 6)

- The stereo imager baseline length (dX):
- Surface plane with respect to left camera
- Imager focal length (f)
- Correlation matching error
- The size of measured surface patch
- The number of pixels (inliers) on the surface plane

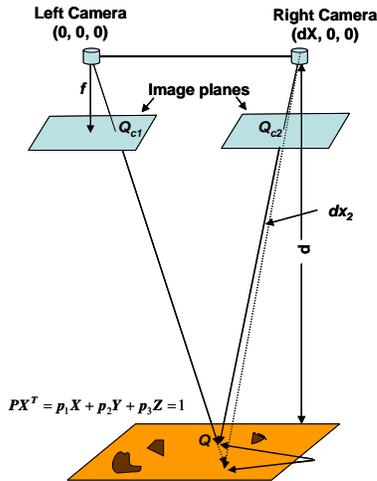


Fig. 6. The stereo-based surface slope estimation model.

Consider an underlying surface P represented by $p_1X + p_2Y + p_3Z = 1$. The left camera is located at $(0, 0, 0)$ and the right camera at $(dX, 0, 0)$. Let a homogenous point Q_{c1} in the left camera plane be $(x_{c1}, y_{c1}, 1)$ and its projection on P given by:

$$Q(rx_{c1} \quad ry_{c1} \quad r) \quad \text{where} \quad r = 1/(p_1x_{c1} + p_2y_{c1} + p_3)$$

The point Q in the right camera frame is $(rx_{c1} - dX \quad ry_{c1} \quad r)$.

Then its projection in the right camera plane (x_{c2}, y_{c2}) is:

$$x_{c2} = (1 - p_1dX_1)x_{c1} - p_2dXy_{c1} - p_3dX \quad (1)$$

$$y_{c2} = y_{c1}$$

Expressing (1) from camera frame to image frame gives:

$$x_2 = (1 - p_1dX)x_1 - p_2dXy_1 + (-p_3f + p_1x_0 + p_2y_0) \quad (2)$$

$$y_2 = y_1$$

where f is the focal length, (x_0, y_0) is the camera center.

We rewrite the first equation in (2) to:

$$x_2 = a_1x_1 + a_2y_1 + a_3 = AX_1 \quad (3)$$

where

$$A = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} -dX & 0 & 0 \\ 0 & -dX & 0 \\ x_0dX & y_0dX & -fdX \end{pmatrix} \begin{pmatrix} p_1 \\ p_2 \\ p_3 \end{pmatrix} = TP$$

If the coefficients (a_1, a_2, a_3) are known, the surface plane (p_1, p_2, p_3) can be easily derived by simple linear manipulations as:

$$P = T^{-1}A \quad (4)$$

Next, assume that each pixel in a patch of left image ($a < x < b$, $c < y < d$) is matched to the right image and that the matching error follows a Gaussian zero-mean distribution. The coefficients (a_1, a_2, a_3) can be estimated by a least square method:

$$F = (AX_1 - X_2)^T COV(X_2)^{-1} (AX_1 - X_2) = \min \quad (5)$$

and

$$\hat{A} = (X_1^T COV(X_2)^{-1} X_1)^{-1} X_1^T COV(X_2)^{-1} X_2$$

where $COV(X_2)^{-1}$ is the inverse of the measurement error covariance matrix.

The covariance matrix of A can be obtained by linear error propagation as:

$$COV(\hat{A}) = (X_1^T COV(X_2)^{-1} X_1)^{-1} \quad (6)$$

Therefore the covariance matrix of P is:

$$COV(\hat{P}) = T^{-1} COV(\hat{A}) (T^{-1})^T \quad (7)$$

The slope of the surface P is:

$$\theta = \arctan\left(\frac{(p_1^2 + p_2^2)^{1/2}}{p_3}\right) \quad (8)$$

Now, let

$$d\theta = \frac{\partial\theta}{\partial p_1} dp_1 + \frac{\partial\theta}{\partial p_2} dp_2 + \frac{\partial\theta}{\partial p_3} dp_3 = (dp_1 \quad dp_2 \quad dp_3) K \quad (9)$$

$$K = \begin{pmatrix} p_1 p_3 \\ f_1 f_2^{1/2} \\ p_2 p_3 \\ f_1 f_2^{1/2} \\ -f_2^{1/2} \\ f_1 \end{pmatrix} \quad \text{where} \quad \begin{aligned} f_1 &= p_1^2 + p_2^2 + p_3^2 \\ f_2 &= p_1^2 + p_2^2 \end{aligned}$$

The covariance of slope θ is then:

$$\begin{aligned} COV(\hat{\theta}) &= \mathbf{K}^T (dp_1 \ dp_2 \ dp_3)^T (dp_1 \ dp_2 \ dp_3) \mathbf{K} \quad (10) \\ &= \mathbf{K}^T COV(\hat{P}) \mathbf{K} \end{aligned}$$

The 1σ slope error plots for the 10m to 100m span, illustrated in Fig. 7, show that for the wall surface (5.6 m by 2.8 m) the slope error is smaller than 3° for 1σ at 100 m. The analytical model described predicts that for a wall twice as large, the slope error would be smaller than 1.5° at 100m.

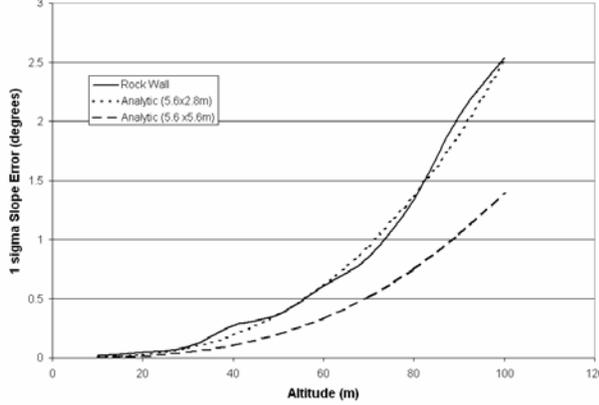


Fig. 7. The surface slope estimation errors comparison between the analytic model and experimental study.

B. Stereo-based Rock Detection

In this section, we develop a model for rock detection and false alarm probabilities for the rock wall data set specifically and compare the model to experimental results; in section V, we extend this to an overall model for the probability of a successful landing given a more general distribution of rock sizes at the landing site.

Our detection and false alarm models for rocks are based on simple, Gaussian models of uncertainty in estimated rock heights above the nominal ground surface. To derive these, we introduce two new parameters for HD, illustrated in Fig. 8. The first denotes the lander rock tolerance T . A rock taller than T could cause a mission failure. The second parameter is the HD algorithm threshold t , used to decide whether or not a detected rock is a hazard. Given the uncertainty in rock height estimation, t is set below T to minimize missed detection of true hazards at the expense of an increased false alarm rate. An appropriate setting for t is then that which minimizes the probability of mission failure (Section V).

The height uncertainty model treats rock height estimation as zero mean Gaussian with range uncertainty (σ). In theory:

$$\sigma = \frac{d^2 \cdot ivFOV \cdot k}{dX} \quad (11)$$

where d is the range in meters to the ground surface, $ivFOV$ is the angular resolution of the sensor in mrad. k is the pixel precision, and dX is the stereo baseline in meters.

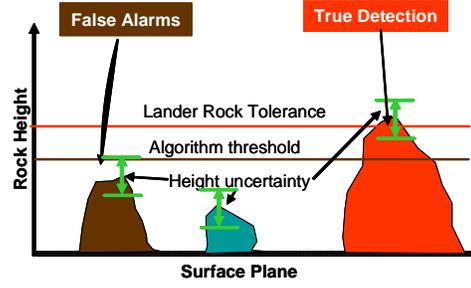


Fig. 8. The lander rock tolerance threshold T (red line) and algorithm threshold t (black line).

The model uses a fixed number of range points on the ground surface as clutter, also represented by a zero mean Gaussian. The probability of detection, P_d , and the false alarm rate (FAR) are computed by integrating the tails of the Gaussian distributions.

$$P_d = \left(\sum_{i=1}^n \int_t^\infty g(x, \sigma_r, H_i) dx \right) / n \quad H_i > T \quad (12)$$

$$FAR = \sum_{j=1}^m \int_T^\infty g(x, \sigma_r, H_j) dx + k \int_T^\infty g(x, \sigma, 0) dx \quad H_j < T$$

where H is the rock height .

The experimental results using the wall data set and the analytic model are shown in Fig. 9. Note that below 60 meters “altitude” stereo-based HD has almost perfect detection.

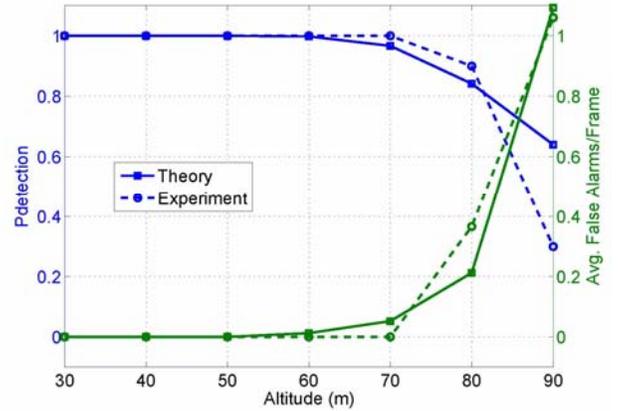


Fig. 9. The hazard detection and false alarm rate comparison between analytic and experimental results of the wall data set.

C. Shadow-Based Rock Detection

We have tested shadow-based rock detection with aerial images of a rock field on Mars Hill in Death Valley, California. A small portion of one such image was shown earlier in Fig. 3. The dataset does not have rock height ground truth, but it includes seven different sun incidence angles between 30° and 70° off nadir. We manually registered these images and manually constructed ground truth of rock footprints by

outlining rocks in the imagery. We then computed the average detection and false alarm rates for 136 reference rocks with diameters greater than 5 pixels in the imagery over the seven sun incidence images of the same location shown in Fig. 5. The plots shown in Fig. 10 summarize the results. The overall probability of detecting rocks with diameters ≥ 5 pixels was 85% with an average of 3 false alarms per image. Perfect performance (100% detection with no false alarms) was achieved for rocks ≥ 25 pixels in diameter. These results can be used to choose the camera field of view and operating altitude to achieve a desired level of reliability in shadow-based rock detection. For example, if hazardous rocks have diameters of 1 m or more, operation at 200 m altitude is desired, and a performance equivalent to the 25 pixel diameter case in Fig. 9 is desired, the camera angular resolution must be 0.2 milliradians/pixel. Since true rock height was not available with this data set, we did not try to evaluate height estimation.

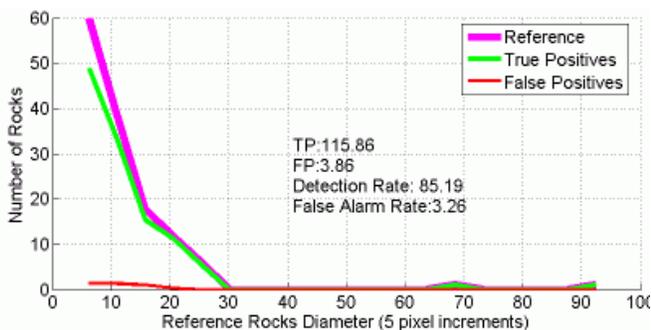


Fig. 10. Shadow-based rock detection. Average detection for seven different sun angles. The minimum number of pixels in a shadow region is 5.

A number of experiments were conducted using the wall dataset. Fig. 11 shows a result for a simulated wall image at 400 m distance. The ground sampling distance (GSD) for this image is 11.2 cm and the RMSE of the difference between the measured height (magenta bars) and actual rock heights (cyan bars) is 1.8 cm, i.e. 5.4% of the average true height (33 cm) of the rocks.

These experiments suggest that error in height estimation grows nearly linearly with altitude. Rock heights can be compensated for sun aureole angle which explains in part the consistent height underestimates. The angle subtended by the Sun from the Earth's point of view is about 0.5 degrees and the forward scattering by the Sun's aureole [8] and can result (on Earth) in a total subtend angle of up to 5° [9]. The net effect is a penumbra along the shadow boundaries that has a width proportional to the rock height. The penumbra represents a transition zone from direct to diffuse illumination resulting in under segmented shadow regions, thus the underestimate in rock height.

The Mars Reconnaissance Orbiter (MRO) entered Mars orbit in the Fall of 2006. Among its instruments it carries the High Resolution Image Science Experiment (HiRISE) instrument. From an altitude of 300 km and with an FOV of $1.14^\circ \times 0.18^\circ$ it is capable of acquiring image swaths 20,264 pixels across having a resolution (GSD) of 30 cm/pixel. The

swath length is typically twice the swath width, thus covering an area 6.2 km x 12.4 km. The instrument was targeted from December 2006 to March of 2007 to acquire high resolution images overlapping three potential landing sites for the Phoenix Mission, scheduled for launch in August of 2007. The 36 GByte dataset covers an area of approximately 1,500 km² in 46 HiRISE images at ~ 30 cm/pixel GSD. At this resolution and given the favorable sun elevation angles (approx. 60° incidence,) the shadow-based rock detection algorithm was applied to map large rocks, on the order of 1 m or larger.



Fig. 11. Shadow-based rock detection from wall dataset simulated to 400m altitude. The bar chart compares true (cyan) to estimated (magenta) rock height. The error in height estimation grows nearly linearly with altitude.

A rock mapping algorithm, derived from the real-time HD algorithm described above affords the additional computational cost needed for a more refined analysis of the shadows detected at the available GSD, and the processing of the very large images in the dataset. These modifications include analyzing the detected shadow regions in more detail to attempt to separate merged shadows from adjacent or close rocks. A second refinement analyses the terminator (the illumination boundary between the rock and its self shadow.) Certain terrain features in the scale of the lander cast shadows having lengths comparable to those of rocks. Current risk assessment methods are primarily concerned with rock hazards rather than general hazards. Many of these features are elongated and can be discerned from the aspect ratio of the shadows they cast. Fragmented features however do occur as well, where the fragments have sizes similar to large rocks. The analysis of the illumination gradient along the illumination terminators has been highly successful in discriminating these from large rocks. The mapping algorithm is able to generate a rock description record (position and size of individual shadows and rock models) for an entire HiRISE image in a few minutes. Overall, over 10 million rocks were detected and mapped. Fig. 12 illustrates a "Box", one of three 150 km x 75 km areas designated for study.

With such large areas under consideration, we expect variations in the terrain. Fig. 13 illustrates five representative terrain types in the dataset. Four of them represent areas away from craters whereas the areas near craters typically have large concentrations of rocks, many of them boulders that can severely impact a lander.

Fig. 14 shows an illustrative result of rock detection from a small portion of a HiRISE imagery. The final set of shadows detected (analyzed shadow regions larger than 5 pixels) are illustrated on the bottom by their approximated ellipses. Rock models giving the position and estimated sizes (diameter and height) are derived from these.

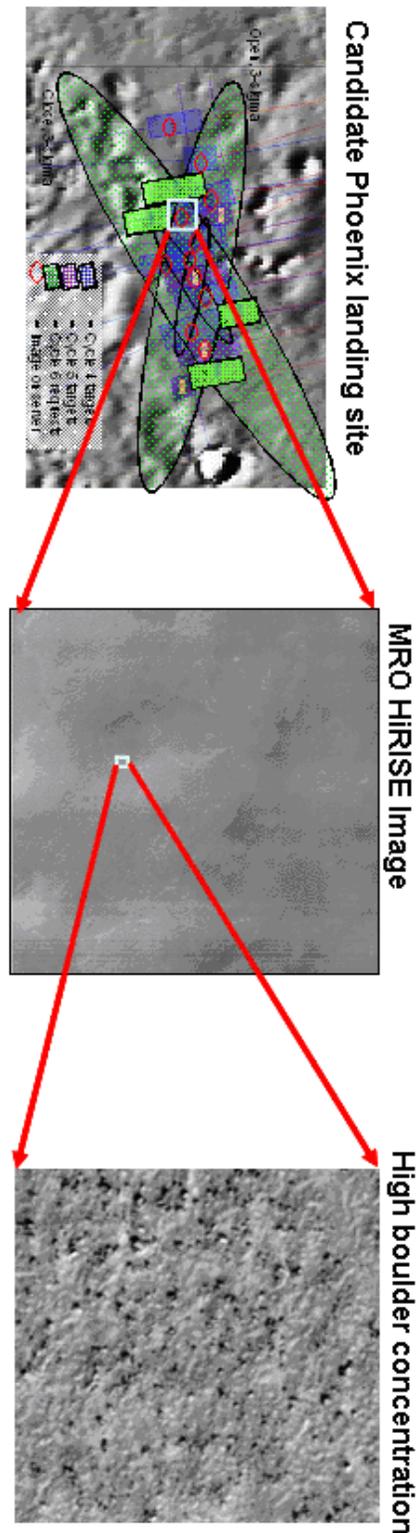


Fig. 12. Potential landing ellipses and HiRISE image coverage for “Box 1”. With a GSD of 30 cm, rocks 1 meter or larger are clearly visible. HiRISE images in the dataset are typically 20,048 pixels across by N pixels long, with N varying from 10,000 to 100,000 pixels. (ADD SCALE markers)

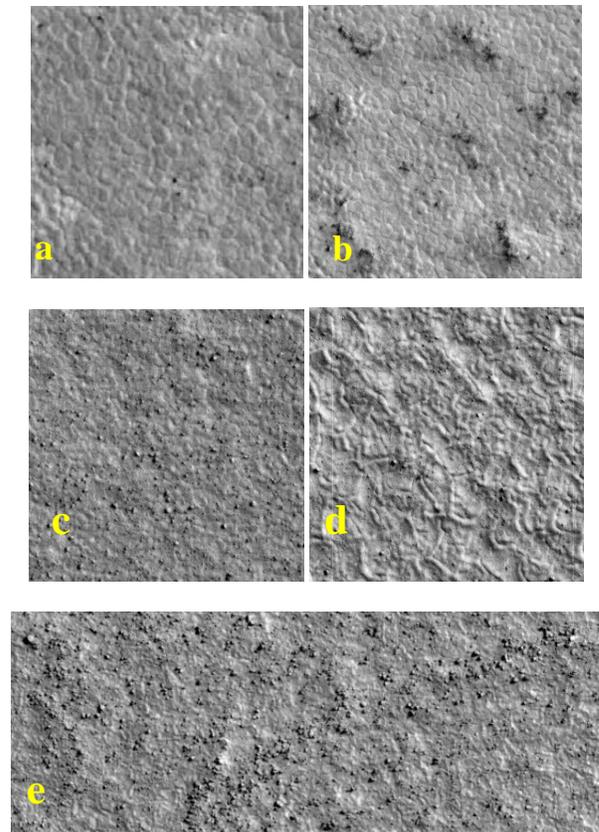


Fig. 13. terrain types in the dataset. (a) Well defined polygons (5-10cm relief at edges) with very few rocks. (b) Boulder clusters interspersed with boulder-free polygonal terrain. (c) Low to medium-density rocks uniformly distributed across surface (d) Rippled terrain with very few rocks or rock-free. (e) Rocks along border of filled-in crater. The boulder on the top left is 4.6 m in diameter

We compared automatic counts with hand counts and with surface counts. These counts are generated by teams led by geophysical scientist that choose and certify landing sites. Hand counts result from manual measuring of less than 200 rocks from small portions of the high-resolution images. Typically two windows are chosen from each image to count and measure rocks in one high- and one low-rock density areas. High rock density windows are typically 300 pixels x 400 pixels (10 m x 1 2m) with about 140 rocks. Low rock density are about 800 pixels x 800 pixels (240 m x 240 m) and contain about 30 rocks. So far, about some 50 such windows have been processed manually. Surface counts are the result of an extended effort to count and map the rocks at the landing sites of previous Mars missions. In particular, the Viking Landers (VL 1 and VL 2) and Mars Pathfinder (MPF.) Such mappings have been instrumental in the derivation of the rock abundance models we described above [11,12]. For these comparisons we used the model coefficients derived form VL2 surface counts derived using monocular and stereo tools applied to the images acquired by the lander sensors. For

automatic mapping we used portions of the HiRISE images that include the landers. Both manual and surface count comparisons have been accepted by the respective geophysical scientist and collaborators that collected that data, and the full set of automatic mapping results together with the statistics and maps derived from them are being used to help in the site selection task. Fig. 16 illustrates one such comparison for the VL2 landing site.

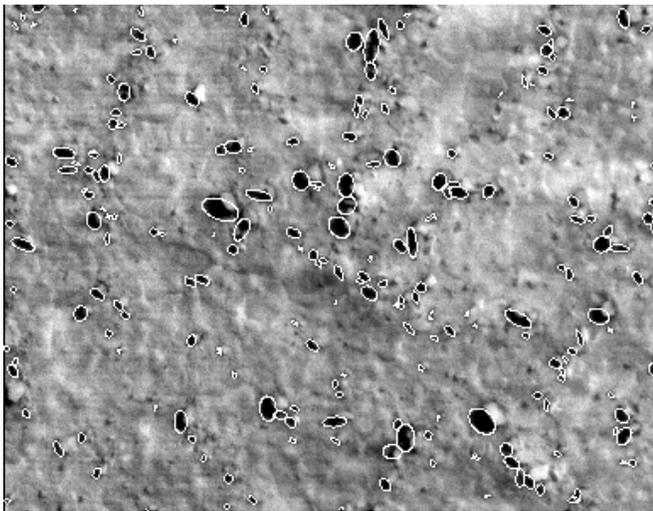
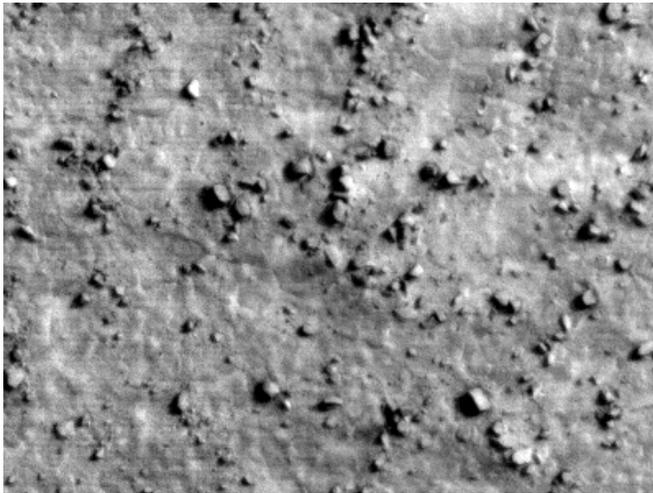


Fig. 14. Detail view of rock field and shadows detected and approximated by ellipses. The smallest rock diameter is about 70 cm.

Fig. 15 illustrates results for an entire HiRISE image covering approx. 6 km across by 18 km. Density and thematic maps are derived from the full rock population in a straight forward manner at any level of granularity. In this example the map cells are 100 m x 100m. The color scale of the density map (center) represent number of rocks having diameters of 1 m or larger. The color scale of the thematic map denotes ranges of rock density related to rock abundance (see caption.)

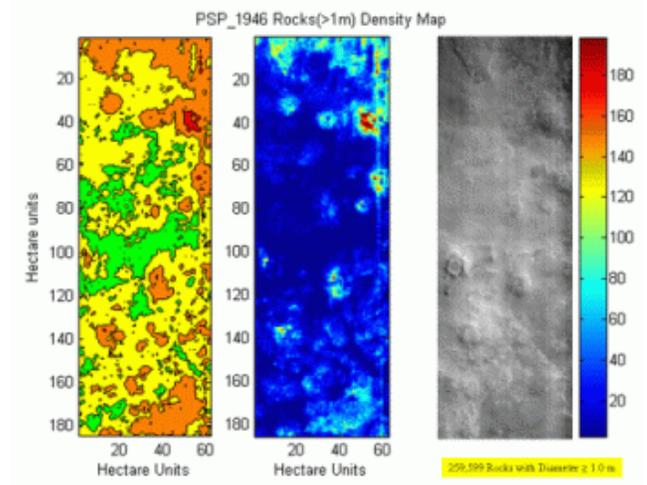


Fig. 15. The HiRISE image (right) covers an area 6 km across by 18 km. The color-coded density map (center) represents number of rocks, larger than 1m, in 100 m x 100m cells. The thematic map (left) encodes density related to rock abundance, and N, the cumulative fractional area (CFA [12]) covered by rocks having a diameter 1 m or larger: Green: ≤ 4 rocks/Ha (~ 5% CFA); Yellow: $4 < N < 32$ rocks/Ha (~10-15% CFA); Orange: $32 < N \leq 128$ rocks/Ha (~10-20% CFA); Red: $N > 128$ rocks /Ha (> 20% CFA)

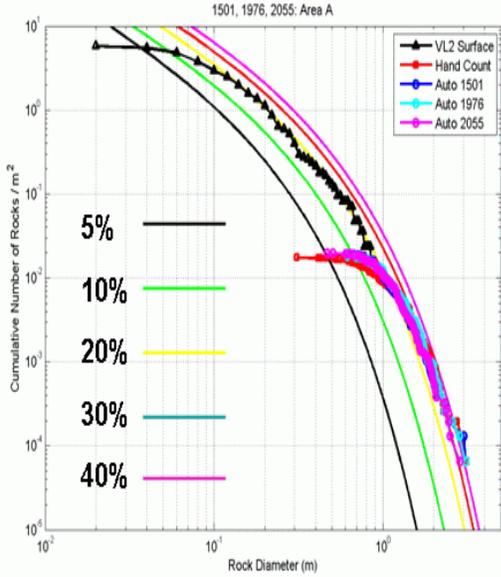
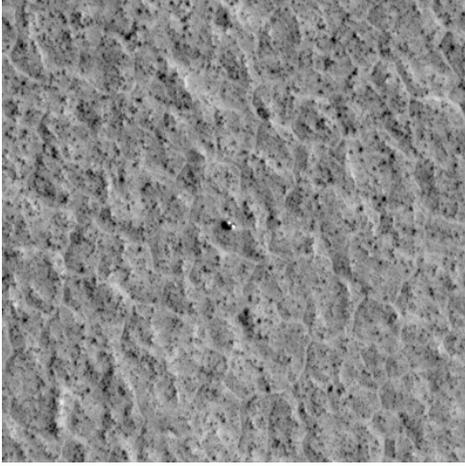


Fig. 16. Validation of automatic rock detection. (top) Viking Lander 2 (at center of image) landing site. The plots represent rock distributions. The surface counts were provided by Matthew Golombek of the Jet Propulsion Laboratory and hand counts by Prof. Raymond Arvidson of the Washington University in St. Louis).

V. SAFE LANDING PROBABILITY MODEL

A great deal of work has been dedicated to Martian rock distributions and selecting safe landing sites on Mars over the last few years [11,12]. The distributions model (also called rock abundance model), is given in terms of the cumulative fractional area covered by rocks as a function of rock diameter. The Martian rock abundance varies from 0 to up to 40%. (Fig. 17) For example, the MSL is targeting to land a rover to the terrain less than 10% rock abundance with landing failure probability due to rock impalement less than 0.25%.

Although we are more interested in rock height, the

statistical studies suggest a relationship between diameter and height [11]. For consistency with these studies we use rock diameter in our models below without loss of generality.

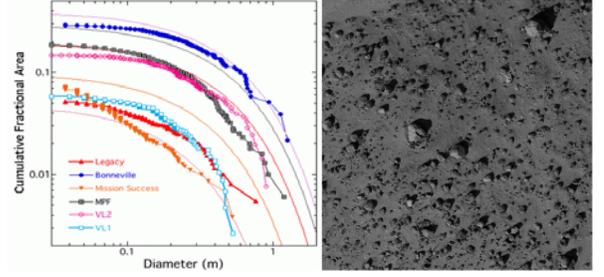


Fig. 17. Rock abundances for selected Mars missions. Mars Pathfinder (MPF) , ~20%; Viking Lander 1 (VL1), 17%; Viking Lander 2 (VL2), 18%; Bonneville, ~32%; Mars Hill (on Earth, image shown), ~22%.

The model we present here is suitable for both for both stereo and shadow based rock detection. Although their errors are influenced by different factors, they are assumed to follow Gaussian distributions. For example, the shadow-based rock detection error varies from 10%, mostly dependent on sun aureole effects, to 50% for the worst case rock shape, i.e. a hemispherical rock. The range is valid for Sun incidence angles between 30 and 70 degrees, our target range.

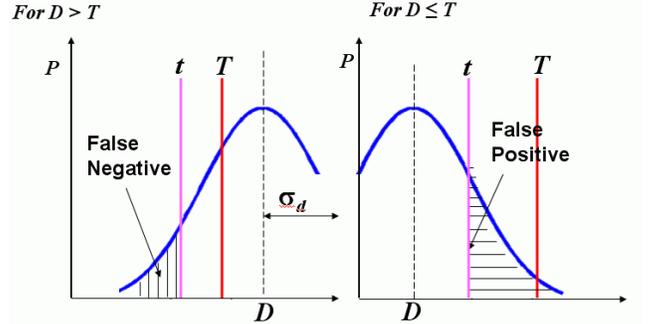


Fig. 18 Uncertainty in the diameter (or height) of a detected rock is modeled by a Gaussian distribution, T represents the mechanical hazard threshold or rock tolerance and t represents the algorithm hazard threshold.

Referring to Fig. 18, the probability of FN (false negatives or, missing a hazard rock) and the probability of FP (false positives, or misdetecting a non-hazardous rock as a hazard) are given by:

$$P_{FN}(D, t) = \int_0^t g(x, \sigma_d) dx \quad \text{for } D > T \quad (13)$$

$$P_{FP}(D, t) = \int_t^\infty g(x, \sigma_d) dx \quad \text{for } D \leq T \quad (14)$$

where, D represents rock diameter. Next we apply these uncertainty models to models of rock distributions. Golombek, et. al. [12] defines a rock distribution F as a cumulative fractional area distribution of rock diameters for a given abundance:

$$F(\mathbf{k}, D) = \mathbf{k} \cdot \exp(-q(\mathbf{k}) \cdot D) \quad (15)$$

with:

$$q(\mathbf{k}) = c_a + c_b / \mathbf{k}$$

\mathbf{k} : abundance; D : diameter

Where C_a and C_b are constants and \mathbf{k} represents rock abundance, the percent of the area covered by rocks. In our model we prefer to describe the rock populations instead by a distribution of the number of rocks of diameter D per square meter. We have:

$$\begin{aligned} f(\mathbf{k}, D) &= \frac{d(1 - F(\mathbf{k}, D))}{dD} / (0.25 \pi D^2) \\ &= \mathbf{k} \cdot q(\mathbf{k}) \cdot \exp(-q(\mathbf{k}) \cdot D) / (0.25 \pi D^2) \end{aligned} \quad (16)$$

The expected numbers of missed (FN) and misdetected (FP) rock hazards, per unit area, are respectively given by:

$$N_{FN}(\mathbf{k}, t, T) = \int_T^\infty P_{FN}(D, t) \cdot f(\mathbf{k}, D) dD \quad (17)$$

$$N_{FP}(\mathbf{k}, t, T) = \int_0^T P_{FP}(D, t) \cdot f(\mathbf{k}, D) dD \quad (18)$$

The expected number of true hazard rocks (HZ) per unit area is:

$$N_{HZ}(\mathbf{k}, T) = \int_T^\infty f(\mathbf{k}, D) dD \quad (19)$$

The expected number of detected hazard (DH) rocks per unit area is then:

$$N_{DH}(\mathbf{k}, t, T) = N_{FP}(\mathbf{k}, t, T) + N_{HZ}(\mathbf{k}, T) - N_{FN}(\mathbf{k}, t, T) \quad (20)$$

Let us now consider the probability of mission failure from a blind landing, i.e., a landing without a hazard detection capability. For that we follow the suggestion of [10] and consider that the number of rocks per unit area is modeled by a Poisson distribution. The probability of exactly n rocks in any given area is given by:

$$P(n, \lambda) = \frac{\lambda^n \cdot \exp(-\lambda)}{n!} \quad (21)$$

$$\lambda = \mu = N_{HZ}(\mathbf{k}, T) \cdot A_L$$

With A_L denoting the area of the lander, λ then represents the expected number of hazard rocks under the lander.

During a blind landing, the probability of failure (p_f), i.e. that at least one rock of a given size, or larger, is within the area A_L is:

$$\begin{aligned} P_f(\mathbf{k}, T) &= 1 - P(0, A_L \cdot N_{HZ}(\mathbf{k}, T)) \\ &= 1 - \exp(-A_L \cdot N_{HZ}(\mathbf{k}, T)) \end{aligned} \quad (22)$$

Now, the probability of mission failure (MF) with *perfect* hazard detection, and the availability of m non-overlapping

landing sites is equivalent to the probability of all m sites having at least one hazard rock:

$$P_{MF}(\mathbf{k}, T, m) = (P_f(\mathbf{k}, T))^m \quad (23)$$

The probability of successfully finding at least one site free from rock hazards would be:

$$P_s(\mathbf{k}, T, m) = 1 - (P_f(\mathbf{k}, T))^m \quad (24)$$

Now, using (17), the probability of not landing on a false negative hazard is:

$$P_{safe}(\mathbf{k}, t, T) = \exp(-A_L N_{FN}(\mathbf{k}, t, T)) \quad (25)$$

The probability of finding a safe site, with m non-overlapping sites available to divert to, is equivalent to the probability that at least one of them is free from any detected hazard. Using (20) we have that:

$$\begin{aligned} P_{safesite}(\mathbf{k}, t, T, m) &= \\ &= 1 - (1 - \exp(-A_L \cdot (N_{DH}(\mathbf{k}, t, T))))^m \end{aligned} \quad (26)$$

The probability of successful HD landing is then:

$$P_{success}(\mathbf{k}, t, T, m) = P_{safe}(\mathbf{k}, t, T) \cdot P_{safesite}(\mathbf{k}, t, T, m) \quad (27)$$

The analytical model of probability of a successful landing provides a tool not only to estimate such probabilities but also to compare blind landing to landing with hazard detection vision capabilities. Fig. 19 shows plots of such comparison instantiated for MSL mission parameters, i.e., a rock tolerance, or mechanical threshold T , of 60 cm, and a 4 m² lander undercarriage. The stereo HD plot (green) is for a sensing altitude of 70m. Note that this prediction is consistent with the wall results at 70m observed in the error propagation model described in Section IV (see also Fig. 9) for the wall dataset. The safe landing probability predictions applied to shadow-based detection are also illustrated in Fig. 19. The blue plot corresponds to the 5.4% height errors from the wall dataset at a simulated 400 m altitude (illustrated above in Fig. 11) The red plot corresponds to a worst-case shape idealized hemispherical rocks illuminated by a 50° Sun incidence angle, and assumes that 30% of the rocks are pyramidal, i.e., not detected because they do not cast shadows. Note however the doubling of rock abundance for a given level of safety. The MSL probability goal is at 0.9975%.

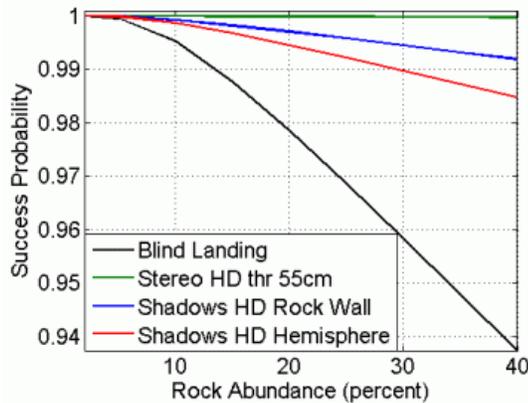


Fig. 19. Landing with HD capability compared to blind landing instantiated for MSL parameters (60 cm rock tolerance, 4 m² lander undercarriage.) The success probability is the product of the probability of no detections and the probability of no missed detections. The green plot predicts stereo-based probability with sensing from 70 m altitude. The blue plot is for the wall simulated 400m altitude result in Fig. 11. The red plot is for idealized hemispherical rocks at 50° incidence illumination.

VI. CONCLUSION

We used a Mars landing scenario as an extreme case of a fast, near-vertical descent to motivate sensor selection for landing hazard detection. This and considerations of minimizing mass, power, and volume while maximizing relevance to other missions led us to conclude that stereo vision and shadow analysis with descent cameras appear to be the smallest sensor suite with the widest applicability, given the state of development of sensor alternatives today. We then outlined algorithms we have developed to date to detect slope hazards with stereo vision and rock hazards with stereo vision and shadow analysis. We derived analytical performance models for these based on Gaussian noise models, compared the prediction of those models to experimental data, and found reasonably good agreement. This implies that the models are useful for predicting performance of these functions in operational scenarios. Therefore, we then embedded the hazard detection performance models in a model for the probability of landing safely, given parameterized models of lander rock tolerance, lander area, and parameterized rock size/frequency distributions fit to Mars and terrestrial data. When this model is instantiated for parameters of the MSL mission, it predicts that even very conservative assumptions about the performance of the vision system will reduce the probability of a failed landing by at least a factor of four compared to a blind landing for any rock abundance. Conversely, for the level of safety desired by MSL, it predicts that the vision system would allow access to roughly triple the fraction of the planet as a blind landing. This would represent a major improvement in access to sites of scientific value for a small increase in sensor payload. Analogous benefits should accrue to missions to other bodies in the solar system.

ACKNOWLEDGMENTS

The research described was carried out at the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the Mars Technology Program under NASA Mars Exploration Program Advanced Technologies NRA *OSS-03-01*) and NASA New Millennium Program (NMP) ST9. The authors are grateful to Dr. Adnan Ansar of JPL and to Dr. Andrew Johnson of JPL for useful discussions on the analytical performance models.

REFERENCES

- [1] Andrew E. Johnson, Allan Klumpp, James Collier and Aron Wolf, "Lidar-based Hazard Avoidance for Safe Landing on Mars," *AIAA Journal of Guidance, Control and Dynamics* 25 (5), October 2002.
- [2] A. Huertas, Y. Cheng, & R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *Proc. IEEE Aerospace Conference*, Big Sky, MO. Mar. 2006.
- [3] L. Matthies, A. Huertas, Y. Cheng and A. Johnson, "Landing Hazard Detection with Stereo Vision and Shadow Analysis," *Proc. Infotech@Aerospace2007 Conference*, Rohnert Park, CA, May 2007.
- [4] J. Vevrka et. al. "Imaging of Small-Scale Features on 433 Eros from NEAR: Evidence for a Complex Regolith," *Science*, Vol. 292. no. 5516, Apr. 2001, pp. 484 – 488.
- [5] P. Christensen, et al "Morphology and Composition of the Surface of Mars: Mars Odyssey THEMIS Result", *Science* 300, 2056 (2003).
- [6] H. Hirschmueller, P. Innocent, and J. Garibaldi, "Real-time correlation-based stereo vision with reduced border errors," *Int'l Journal of Computer Vision*, vol. 47, no. 1-3, pp. 229.246, April-June 2002.
- [7] C. Villalpando, "Acceleration of Stereo Correlation in Verilog", *9th Military and Aerospace Programmable Logic Devices (MAPLD) International Conference*, Washington, D.C., September 2006.
- [8] M. Box and A. Deepack, "Finite Sun Effect on the Interpretation of Solar Aureole," *Applied Optics*, Vol. 20, No. 16, Aug. 1981.
- [9] F. Mims, "Solar Aureoles Caused by Dust, Smoke and Haze," *Applied Optics*, Vol. 42, No. 3, Jan. 2003.
- [10] Harold W. Sorenson "Parameter Estimation", Marcel Dekker, Inc. 1980
- [11] D. Bernard and M. Golombek, "Crater and Rock Hazard Modeling For Mars Landing," *AIAA Space Conference*, Albuquerque, NM Aug. 2001.
- [12] M. Golombek, et. al., "Rock Size-Frequency Distributions on Mars and Implications on Mars Exploration Rover Landing Safety and Operations," *Journal of Geophysical Research*, Vol. 108, No. E12, 2003.